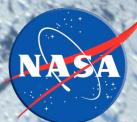


espend eaneige to encitaelloo egred in gninilli eted

Dominic Mazzoni
Machine Learning Systems Group
Exploration Systems Autonomy Section
Jet Propulsion Laboratory

August 24, 2004





Overview

- I. What is Image Data Mining?
- II. What is the Machine Learning Approach to Image Data Mining?
- III. Finding Objects in Images
- IV. Finding Regions in Images
- V. Final Challenge
- VI. Dominic's Recipe





What Is Image Data Mining?

- Have more images than you have time to examine
- Want to answer a question about a phenomenon seen in those images





Examples

- What fraction of the earth's surface is covered by cumuliform or stratiform clouds, by season?
- What is the geographic distribution of the largest volcanoes on Venus?
- What is the mean injection height of smoke plumes caused by forest fires in North America?





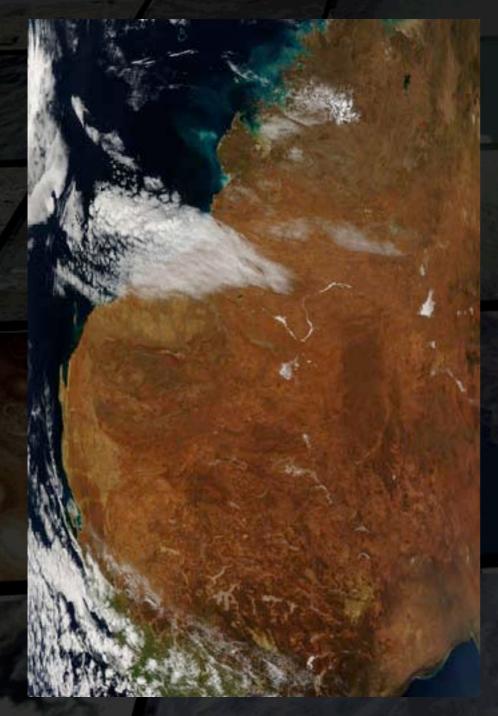
Physics Approach

- Use scientific knowledge about the problem to develop a specific algorithm
- May determine some parameters empirically (e.g. thresholds), but overall approach is grounded in theoretical arguments

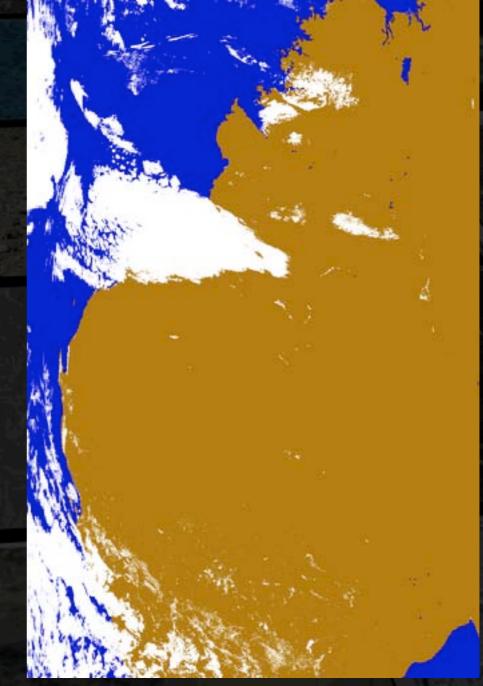




MODIS Cloud Mask







Cloud Mask



MODIS Cloud Mask

19 Spectral bands (of 36)
Sun angle
(from MODIS Level 1)

Snow/ice map from NSIDC
Sea ice concentration from NOAA
(daily)

Land/water mask and elevation data Ecosystem index (constant) Cloud Mask Algorithm

48-bit cloud mask

Overall classification:

Cloudy

Uncertain

Probably Clear

Confident Clear

Other flags (shadow, thin cirrus, suspended dust, etc.)

MODIS: Moderate-Resolution Imaging Spectroradiometer

MODIS Cloud Mask

Band	Wavelength (µm)	Purpose
37 /	0.659	Clouds, shadow
2	0.865	Low clouds
4	0.555	Snow
5	1.240	Snow, shadow
6	1.640	Snow, shadow
7	2.130	Smoke, low clouds vs. snow
17	0.905	Low clouds vs. sun glint
18	0.936	Low clouds vs. sun glint
19	0.940	Shadow
20	3.750	Shadow
22	3.959	Shadow
26	1.375	Thin cirrus, high cloud

Upper tropospheric clouds (below 500 hPa) are the only clouds that significantly affect the 13.9 µm band, thus:

BT(13.9 μ m) < 226 K \Rightarrow high cloud detection

Brightness Temperature

	o temperature		
35	13.935	High clouds	





Stereo Cloud Mask

Nadir Red Radiances

Oblique-angle Red Radiances

Stereo Matching Algorithm

Height Product

SDCM Algorithm



High-conf Cloudy Low-conf Cloudy

Low-conf Surface

High-conf Surface

No match (clear?)

MISR: Multi-angle Imaging Spectroradiometer

(Note: This is a simplification; the full algorithm must also consider winds...)







 It's always important to consider how long the work would take to do manually

help them defeat automatic human detector forms on websites:









Machine Learning Approach

Training Examples

Machine Learning
Algorithm
(Neural Net,
Support Vector Machine,
Naïve Bayes...)

New Examples

Black Box Classifier







JPL's Machine Learning Systems Group



A Word Is Worth...

- Each semantic label ("word")
 increases the computer's
 knowledge substantially more
 than adding data does
- Good training labels are the key to everything!
- A word is worth at least a thousand pixels









Science Goal

- Craters enable us to determine the relative ages of surfaces
- An automated algorithm to detect and measure the size of all craters in a set of images saves time and provides an independent and repeatable estimate



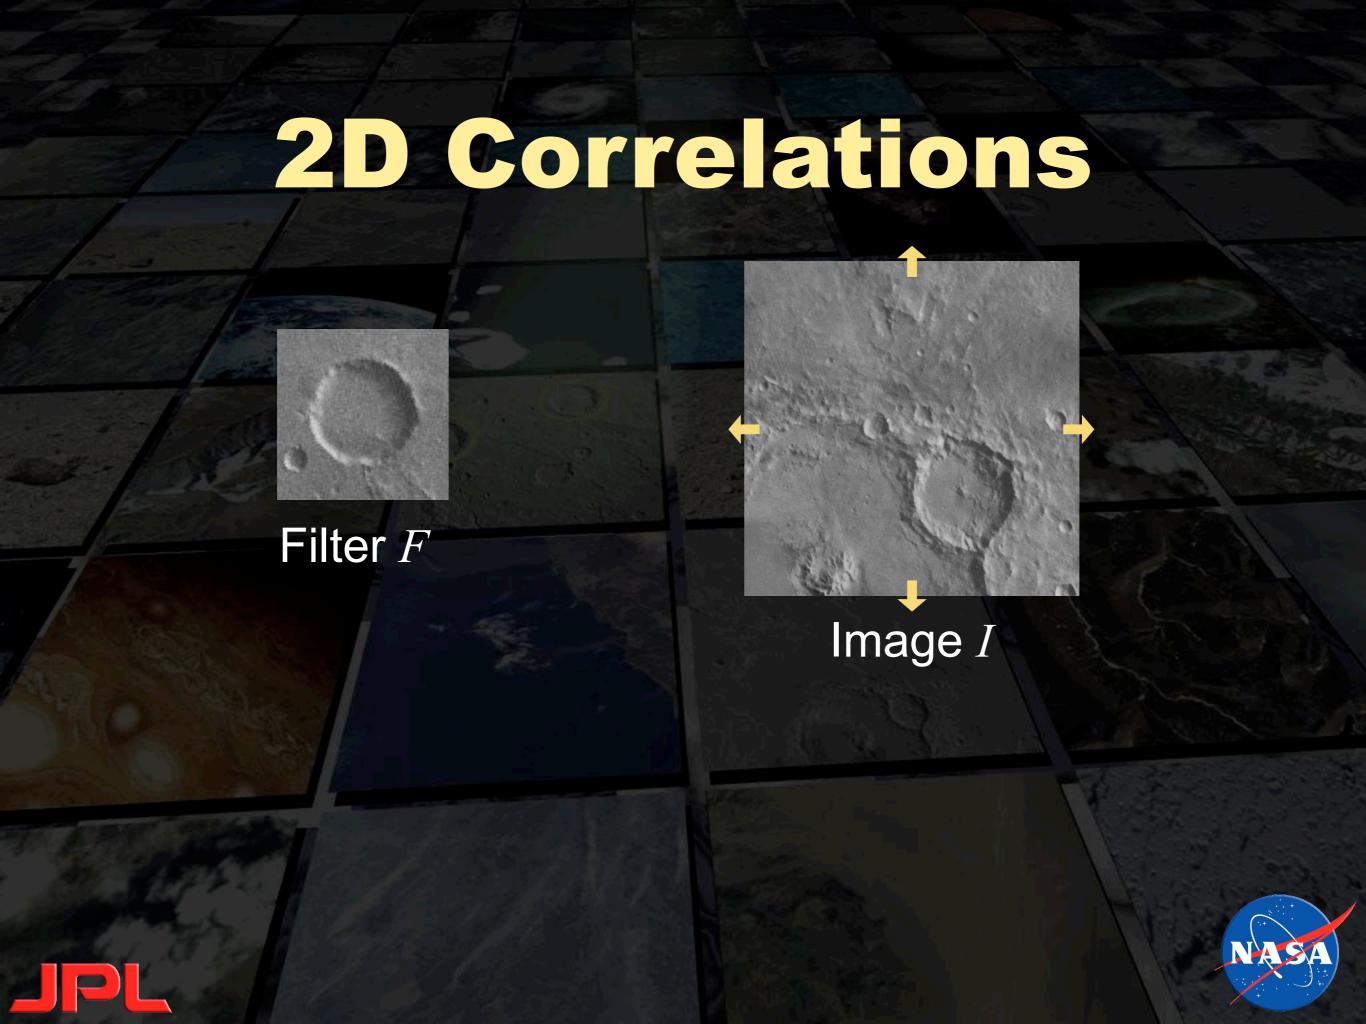




- Pick prototype example(s)
- Build masked filter
- Match
- Arbitrate
- Validate





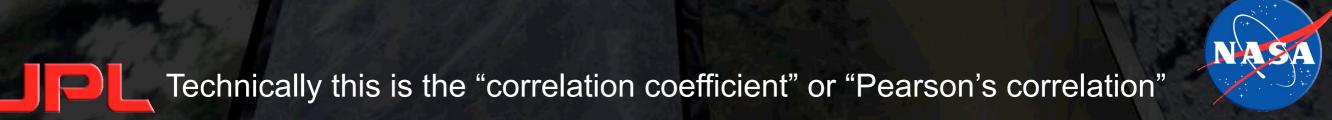


2D Correlations

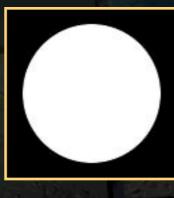
Filter F

Image Chip I'

$$Corr = \frac{E[F \cdot I'] - E[F]E[I']}{\sigma[F]\sigma[I']}$$



2D Correlations



Filter F

Mask M

Image Chip I'

$$Corr = \frac{E[(F \cdot M) \cdot I'] - E[F \cdot M]E[I' \cdot M]}{\sigma[F \cdot M]\sigma[I' \cdot M]}$$





Using Convolution





convolution

$$Corr = \frac{E[(F \cdot M) * I'] - E[F \cdot M]E[I' * M]}{\sigma[F \cdot M]\sigma[I' * M]}$$

$$\sigma^{2}(I * M) = E(I^{2} * M^{2}) - E(I * M)^{2}$$

Note: Don't forget to flip F and M when using convolution...



Convolution using the FFT

- Must zero-pad (F * M) to the size of I.
- Can avoid large FFT computations using an overlap-add technique.





Correlation results



Image I

Filter Response R

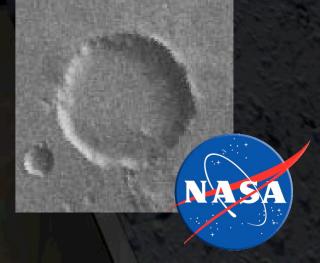
 Arbitration can be used to choose among multiple overlapping matches in the response image.



Rotation and Scale Invariance

- Match objects of different sizes and rotation angles by creating multiple filters
- Can be done efficiently using principal components to compress the set of filters





Incorporating Physical Constraints

When the sun angle is known, sometimes it makes more sense to rotate the image so that the sun is always coming from a particular direction, rather than using rotation invariance





Other applications

- Detecting Volcanoes on Venus from Magellan radar imagery
- Detecting blocks and boulders on the surface of Eros from NEAR
- Detecting lava cones near the Lava Bed National Monument using X-SAR imagery
- Reconnaissance

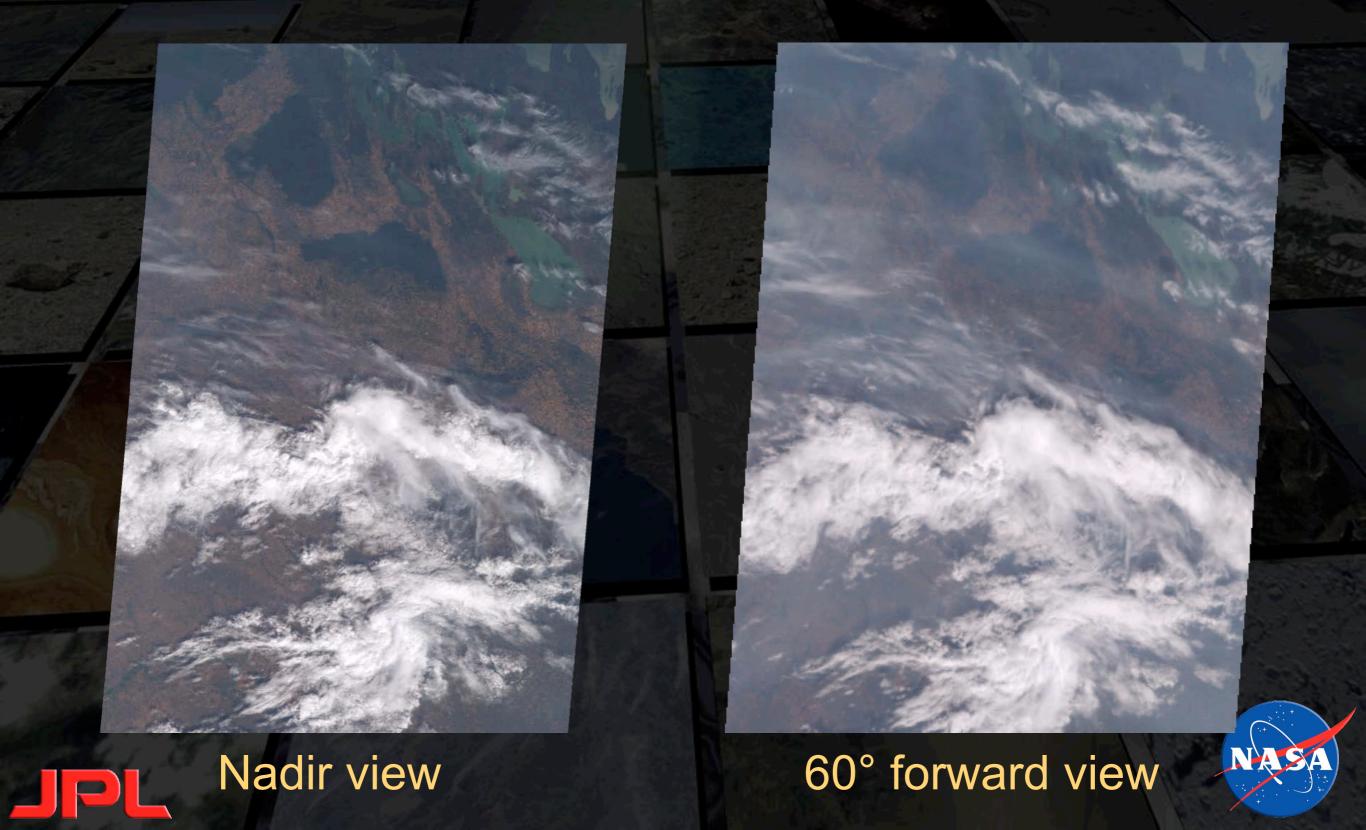








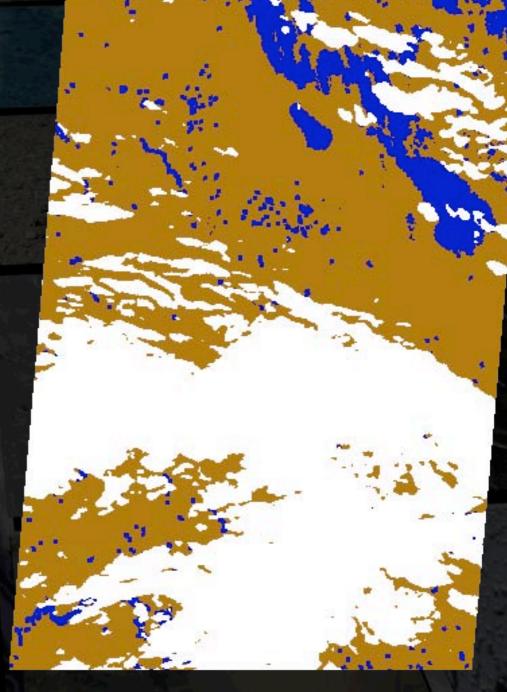
MISR Cloud Mask



MISR Cloud Mask







Cloud Mask



Supervised Training

Training Examples

Machine Learning
Algorithm
(Neural Net,
Support Vector Machine,
Naïve Bayes...)

New Examples

Black Box Classifier











Machine Learning Pitfalls

(This story may be apocryphal.) In the early 1980's, the Pentagon wanted to outfit tanks with digital cameras and advanced computers, allowing them to spot an enemy tank where humans couldn't. They took 100 pictures of tanks hidden in the woods, and 100 pictures of the woods with no tank. Then they carefully trained a neural network on half of the images, and tested the resulting classifier on the other half of the images, achieving remarkably good results!

Unfortunately, it didn't work in practice...it turned out that all of the images with the tank were taken in the morning, while all of the other images were taken in the afternoon. The neural net was looking at much more obvious clues in the image to determine the time of day; it never learned anything about finding an enemy tank.





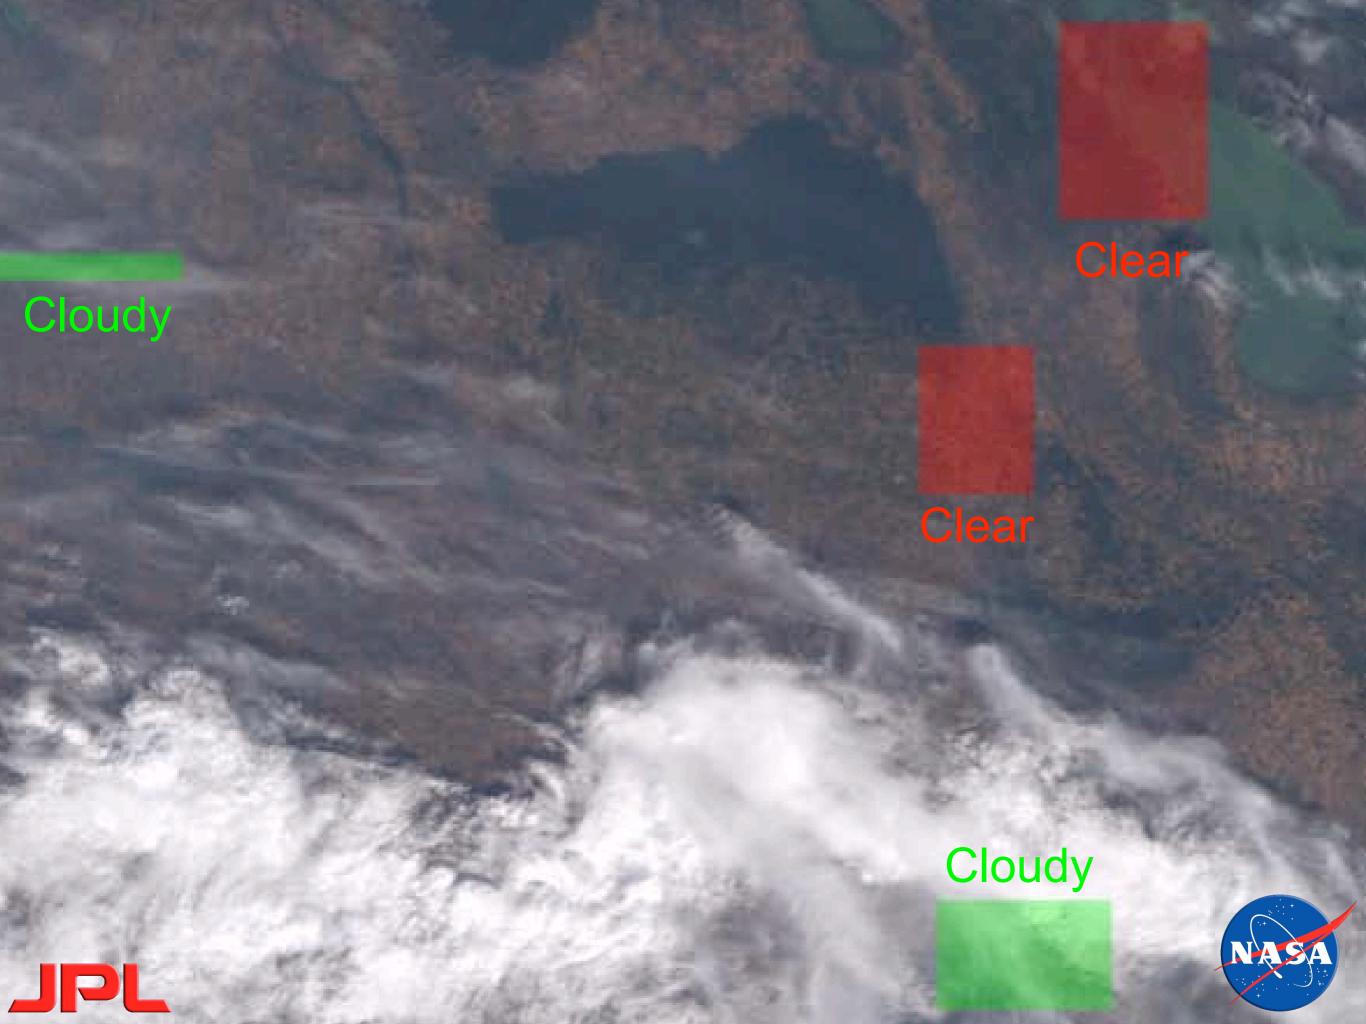
Training a Pixel Classifier

- Label thousands of pixels as positive or negative, ideally in several images
- Choose features that could be used to distinguish between the two classes
- Train a classifier on a fraction of your labels, and validate it on the rest











- Are you defining the problem such that a human expert could answer it?
- To what extent do different human experts agree with each other?





Choosing Features

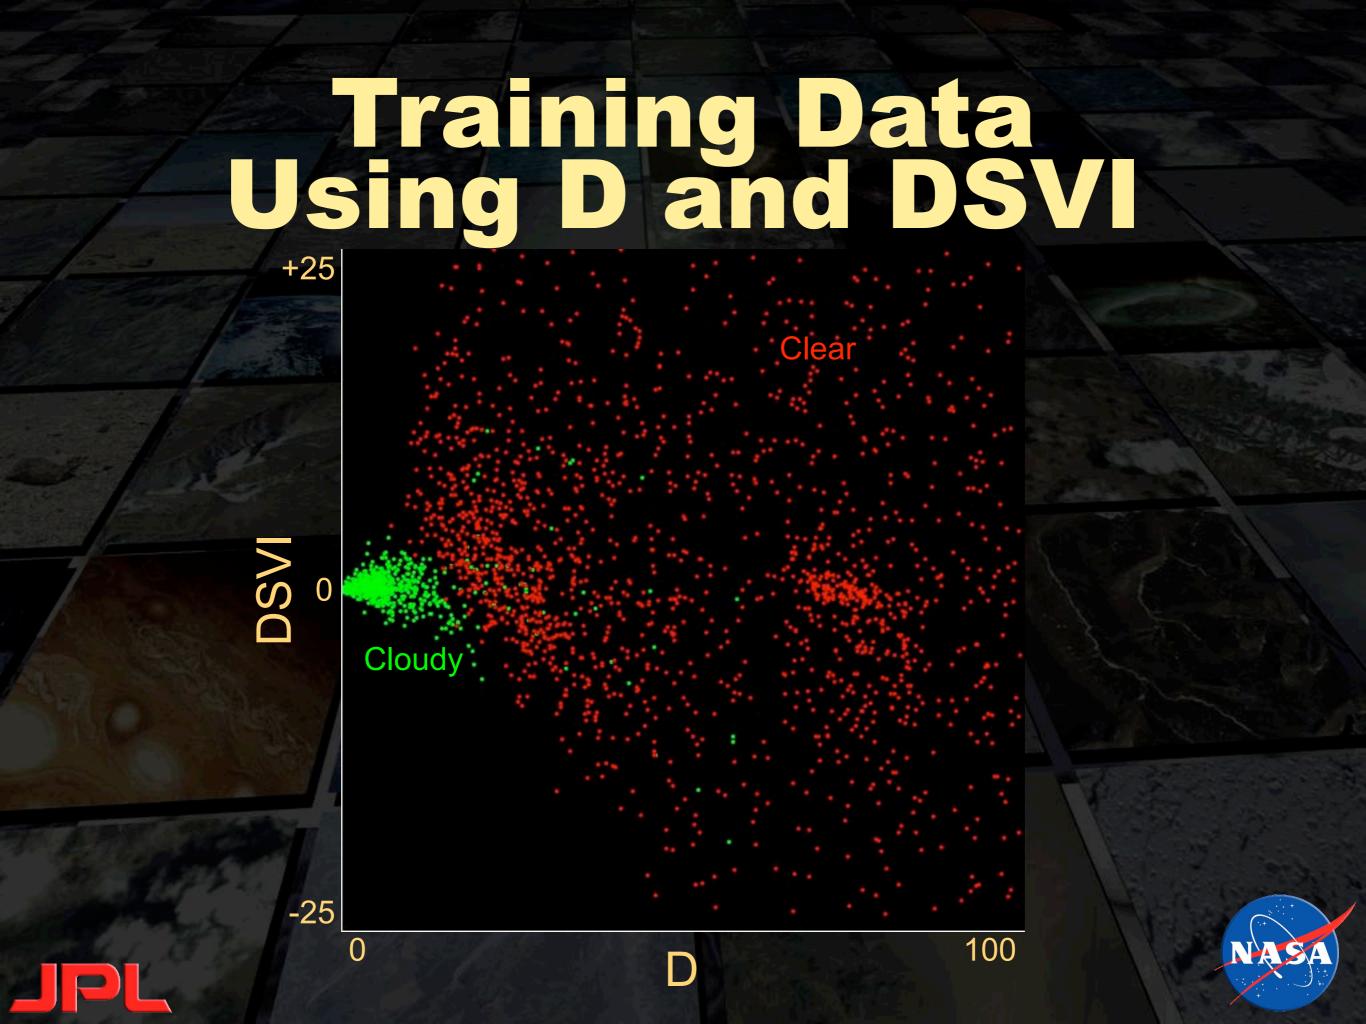
The MISR Radiometric Camera-by-Camera Cloud Mask (RCCM) uses two radiance features for its cloud mask over land:

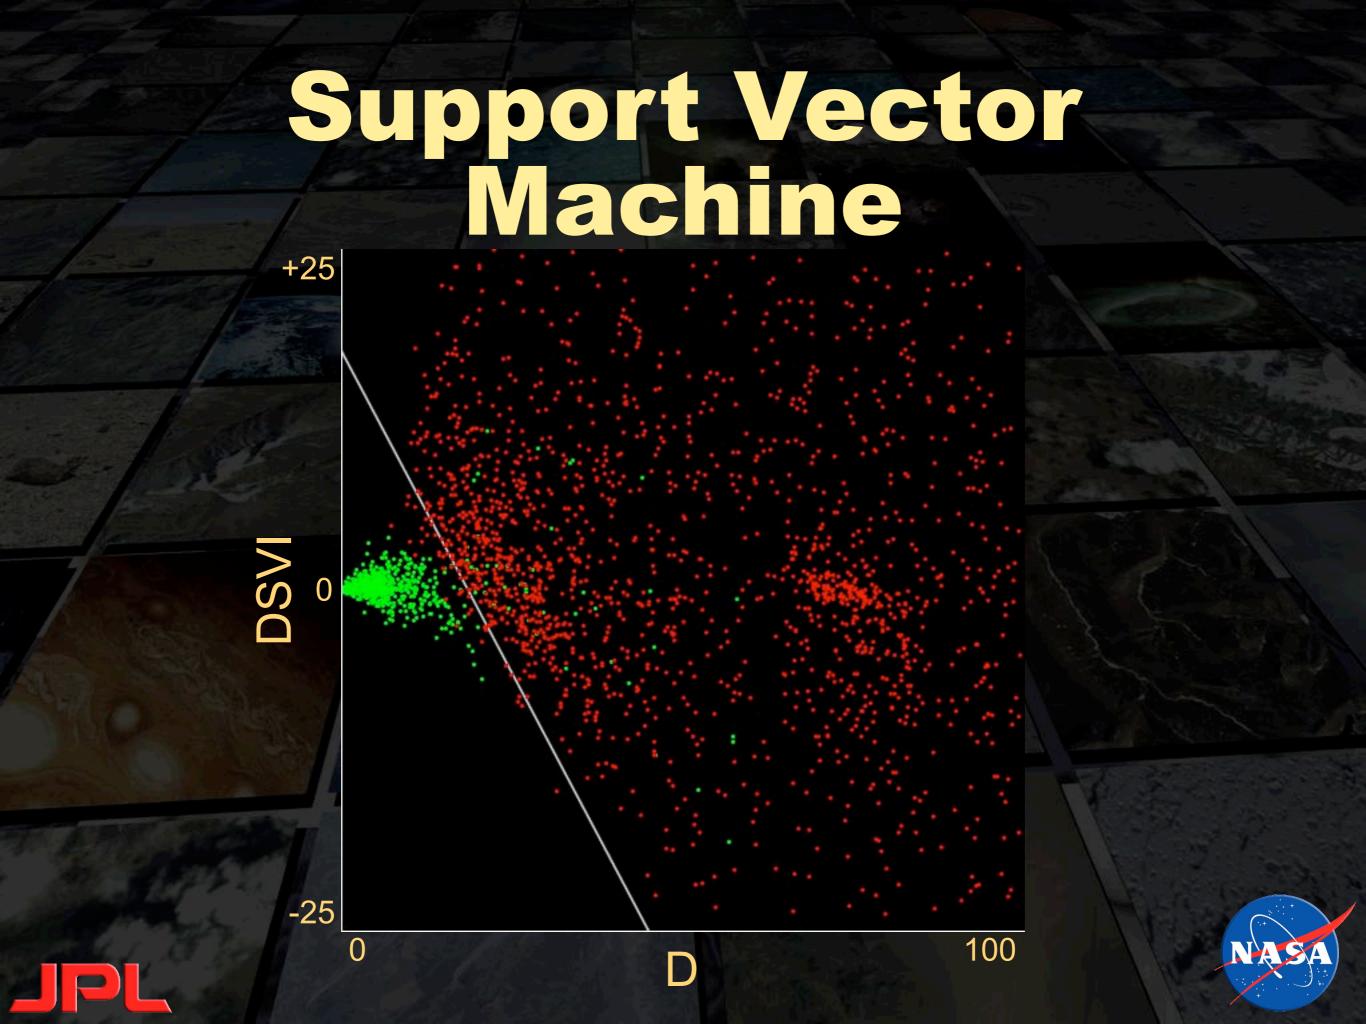
$$NDVI = \frac{Near\text{-}IR - Red}{Near\text{-}IR + Red}$$

$$D = \frac{\left| \text{NDVI} \right|^b}{\text{Red}^2}$$

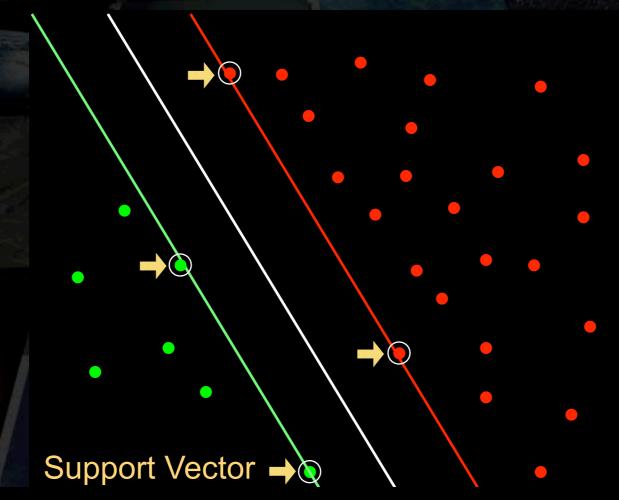
$$b \sim 0.4 \dots 0.6$$







Support Vector Machine



The support vector machine training algorithm finds the separating hyperplane between the positive and negative classes that maximizes the minimum distance between the hyperplane and the nearest points, called *support vectors*.



Nonlinearity

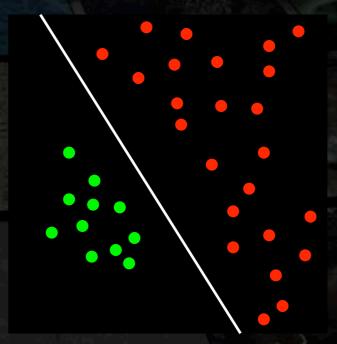
- What if the data is not linearly separable?
 - Develop an algorithm that can find a nonlinear boundary between the classes
 - Or, map the input vectors to a higherdimensional space, including nonlinear features, where the data is linearly separable





Kernels





Feature Space $\Phi(x) \cdot \Phi(y)$

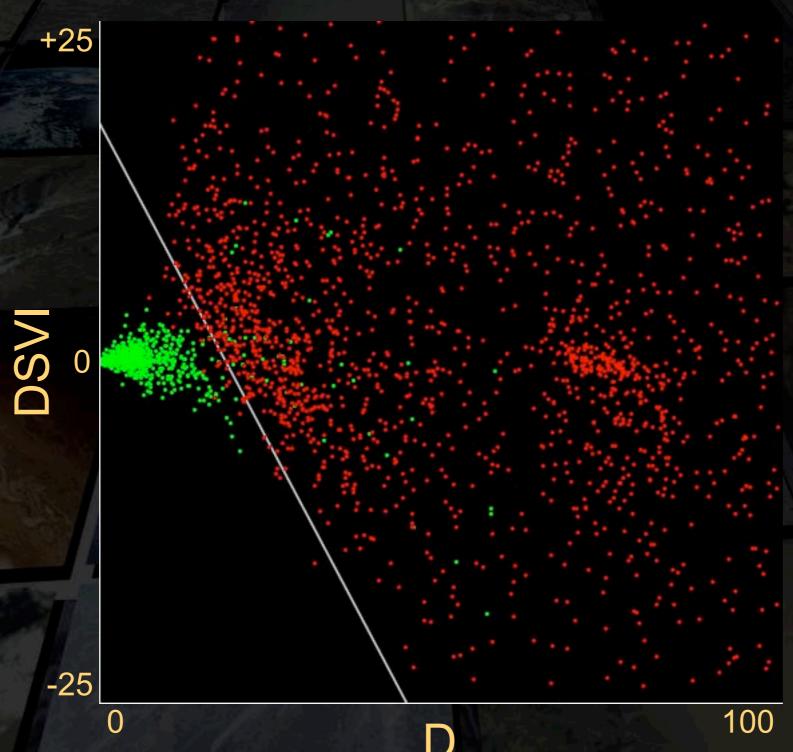
$$K(x,y) = \Phi(x) \cdot \Phi(y)$$





Linear SVM

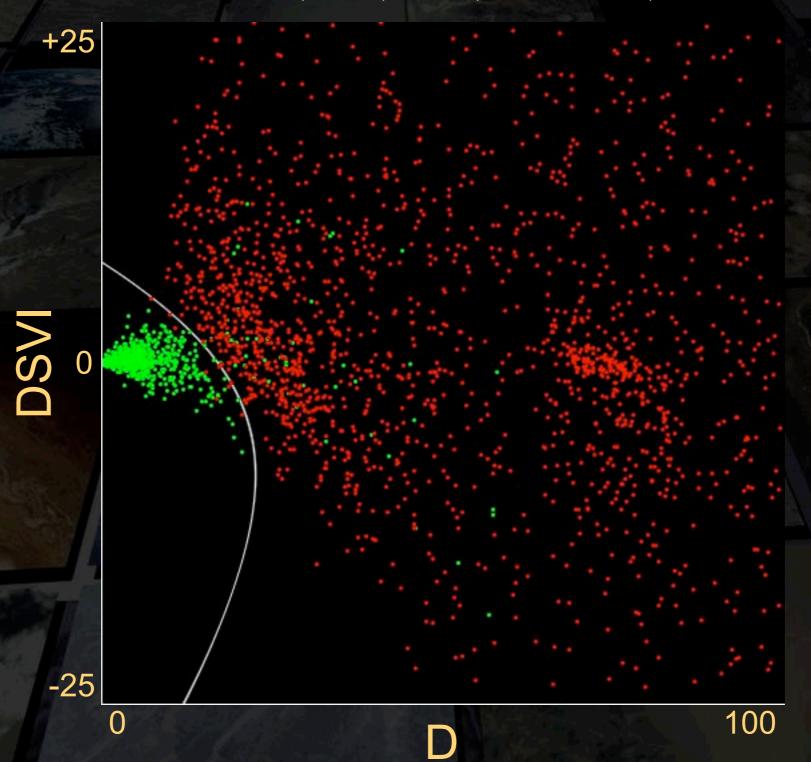
$$K(x,y) = x \cdot y$$





Polynomial Kernel

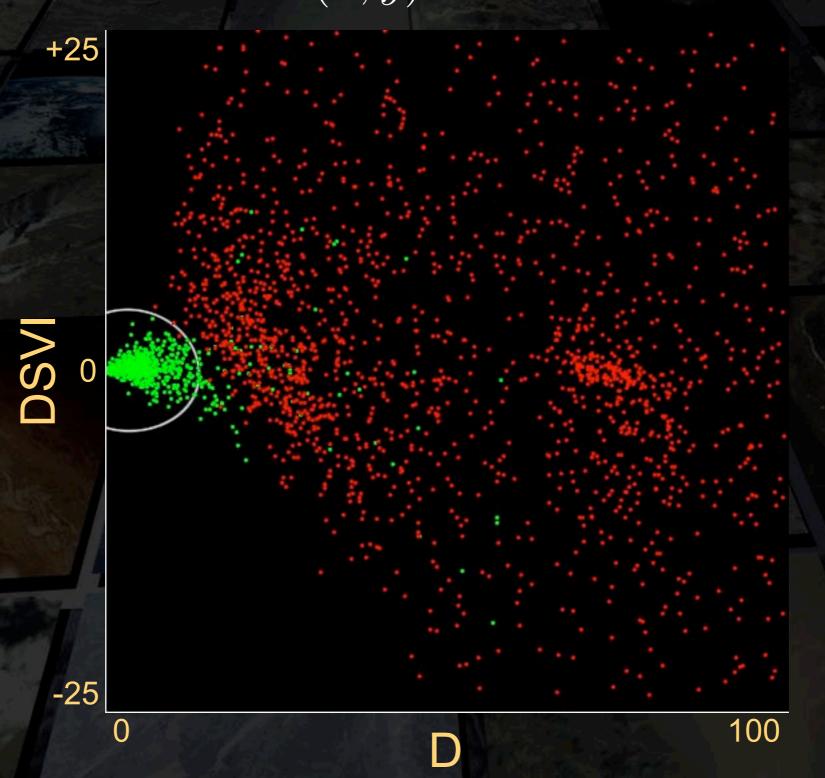
$$K(x,y) = (x \cdot y + b)^p$$





Gaussian Kernel

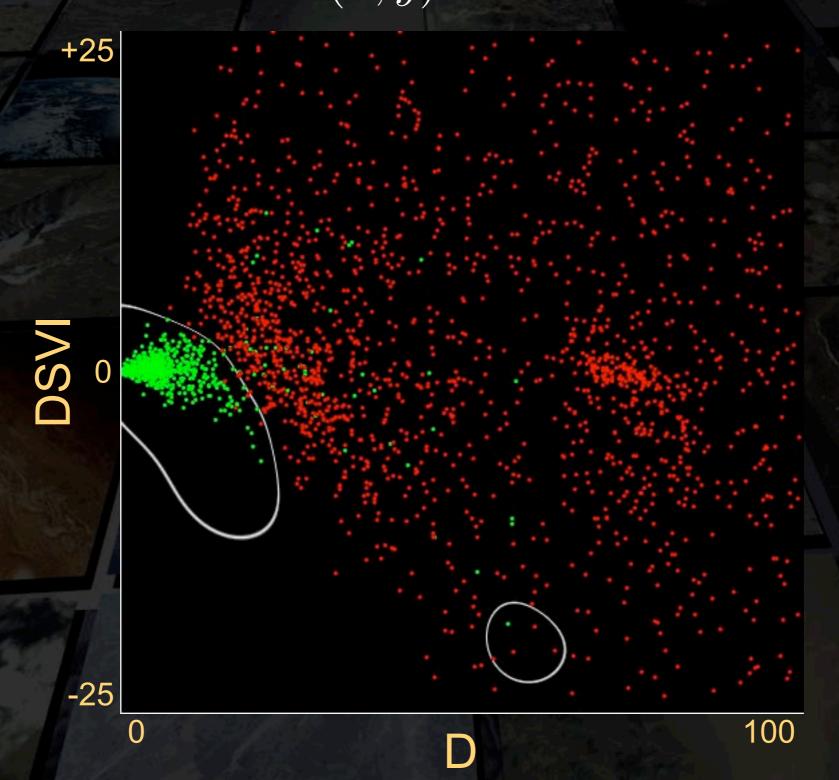
$$K(x,y) = e^{-\gamma|x-y|^2}$$





Gaussian Kernel (2)

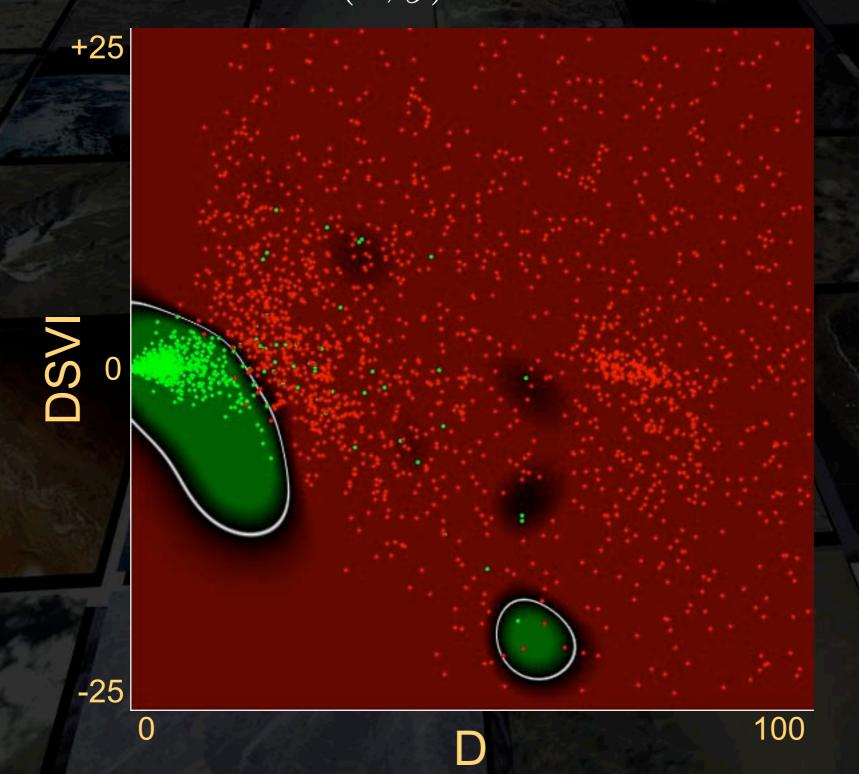
$$K(x,y) = e^{-\gamma|x-y|^2}$$





Gaussian Kernel (2)

$$K(x,y) = e^{-\gamma|x-y|^2}$$

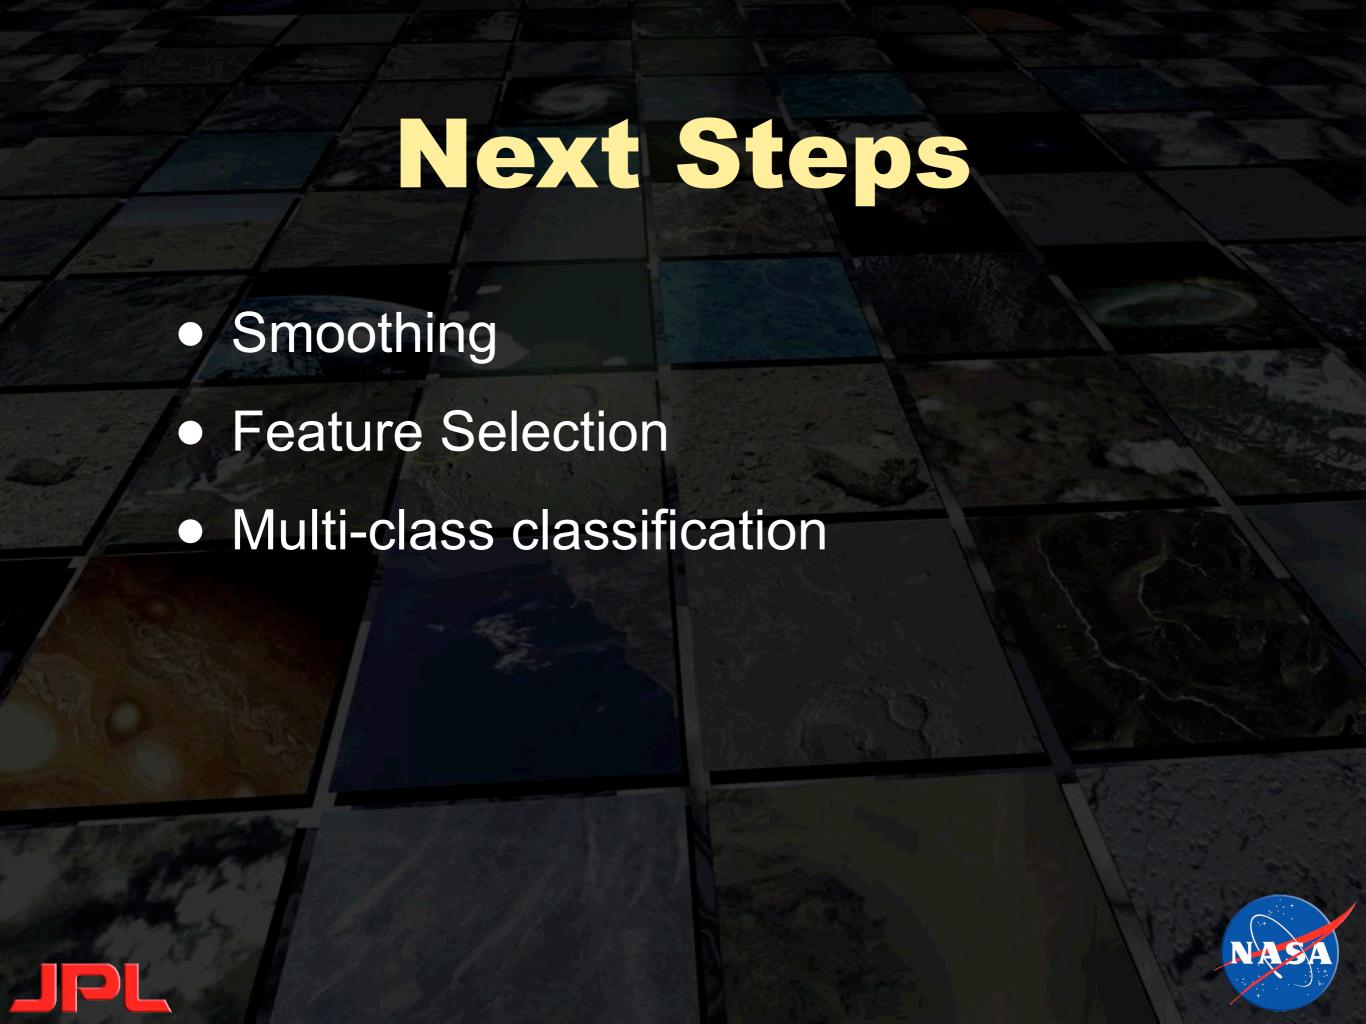




Features	Kernel	Accuracy
D, DSVI	Linear	86.7%
D, DSVI	Poly degree 27	88.9%
D, DSVI	Gaussian	90.7%
3x3 Red	Poly degree 3	91.0%
5x5 Red	Linear	97.6%
5x5 Red	Gaussian	98.0%
5x5 Red, 2 angles	Poly degree 9	98.8%
5x5 3 Colors, 2 angles	Poly degree 9	99.8%







Final Challenge: Smoke Plumes

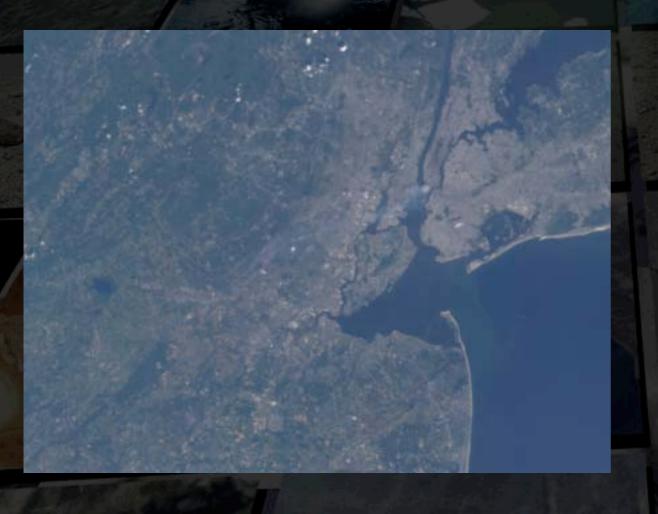
- Automatically detect smoke plumes in multi-angle MISR data
- Identify the plume source, direction, and injection height





Smoke Plumes

New York City, September 12, 2001



Nadir Image

70° Forward Image





Smoke Plumes

Siberia, May 2, 2004

Nadir Image

70° Forward Image



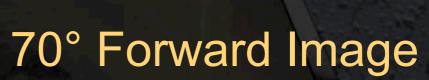


Smoke Plumes

Canada, February 11, 2004



Nadir Image







Approach

- A matching filter ("Finding Objects")
 would not work, because the
 background behind the smoke plumes
 varies too much
- A pixel classifier ("Finding Regions") would not distinguish between smoke plumes, and large areas blanketed by smoke





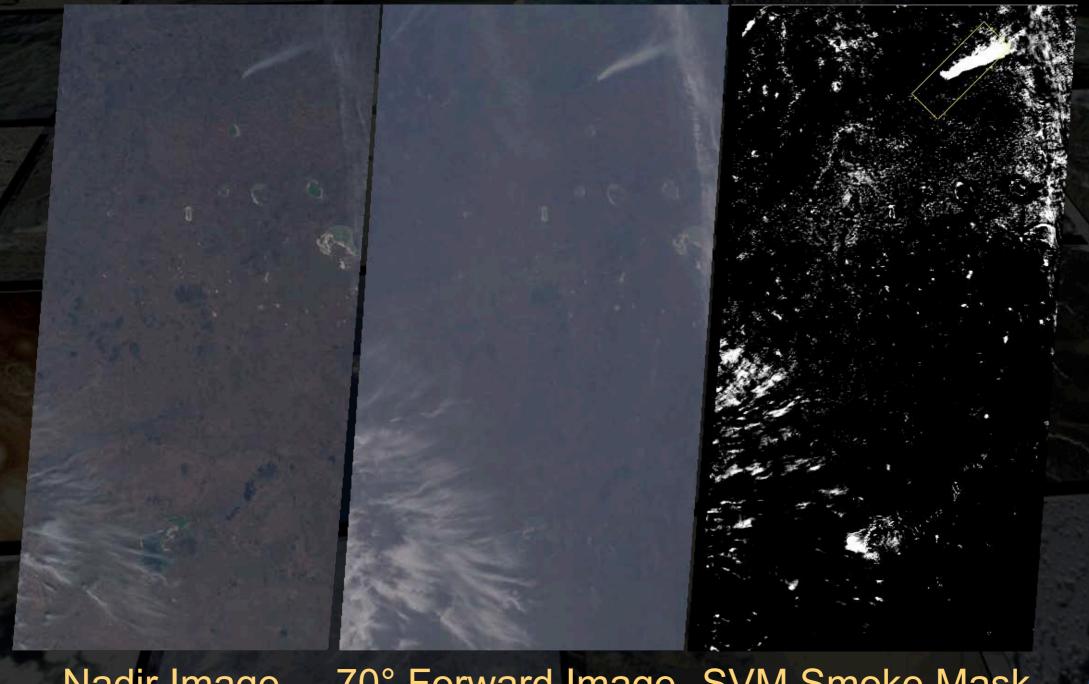
Put them together!

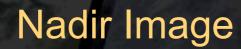
- Use a pixel classifier to identify all of the pixels containing smoke
- Use a matching filter on the resulting classified image to identify smoke plumes





Example Result





70° Forward Image SVM Smoke Mask





Dominic's Recipe

- I. Preheat oven to 475 K
- 2. Start with a well-defined problem
- 3. Decide on approach
- 4. Create training labels





Dominic's Recipe

- 5. Decide on the success criteria
- 6. Use cross-validation to tune the machine learning algorithm on the hand-labeled data
- 7. Run your algorithm on novel data
- 8. Independently validate a fraction of the novel data to determine the final error rate







More Information

- Finding Objects
 - M.C. Burl et al., "Automated Detection of Craters and Other Geological Features", International Symposium on Artificial Intelligence, Robotics and Automation for Space. Montreal, CA, June 2001.
 - M.C. Burl et al., "Automating the Hunt for Volcanoes on Venus", In Proceedings of the 1994 Computer Vision and Pattern Recognition Conference (CVPR-94). Los Alamitos, CA: IEEE Computer Society Press, pp. 302-309.
- Finding Regions / Pixel Classifiers
 - A paper on using Support Vector Machines to develop a MISR Cloud Mask is forthcoming. Please contact Dominic Mazzoni for more information: Dominic.Mazzoni@jpl.nasa.gov
- Support Vector Machines
 - http://kernel-machines.org/
- JPL Machine Learning Systems Group
 - http://ml.jpl.nasa.gov/
 - Supervisor: Rebecca.Castano@jpl.nasa.gov

For other questions, please email: Dominic.Mazzoni@jpl.nasa.gov





Acknowledgements

- MISR research collaborators: Dave Diner, Roger Davies, Mike Garay, Akos Horváth, Amy Braverman
- SVM research collaborators: Benyang Tang, Dennis DeCoste
- Matching filter researchers: Michael Burl, Joe Roden, Tim Stough, Nghia Tang
- Special thanks to Elaine Dobinson for organizing the talk
- Additional thanks: Becky Castaño, Tim Stough, Ben Bornstein, Kiri Wagstaff, Katy Villarreal, Marcia Young
- Background image credit: Dominic Mazzoni and Dave Meyer
- Grad student image credit: Mike Garay
- Caution, tank, corn, handshake, scientist, voodoo doll images all licensed from iStockPhoto
- Dog images credits: Shari and Gary DeGuire
- MISR Data obtained from the NASA Langley Research Center Atmospheric Sciences Data Center; MODIS Data obtained from the Goddard Earth Sciences Data Center; MISR instrument visualization by Shigeru Suzuki and Eric M. De Jong, Solar System Visualization Project.
- Crater images from Mars Mosaicked Digital Image Model CD, 2001 ed., V2
- Primary sponsors: NASA Intelligent Systems Program, National Reconnaissance Office